DOE Center of Excellence Performance Portability Meeting Agenda (effective 4/6/2016)

Day One

Session		Speaker/Topic	Affiliation	Title	Length (m)
	7:30	Registration/Coffee/Light Breakfast		Mingle	0:45
Session chair:		**Overviews**	**Each Center of Excellence to give an overview. Projects, how help is supported, how vendors are integrated, etc…**		
Rob Neely	8:15	Welcome/Kickoff			0:15
	8:30	Tjerk Straatsma	ORNL	Summit COE / CAAR Overview	0:10
	8:40	Jack Deslippe	LBL	NERSC-8 COE / NESAP Overview	0:10
	8:50	Rob Neely	LLNL	Sierra COE Overview	0:10
	9:00	Hai Ah Nam	LANL	Trinity COE Multi-lab Overview	0:15
	9:15	Katherine Riley	ANL	ANL COE Overview	0:10
	9:25	Katherine Riley	ANL	HPCOR Workshop Recap	0:10
	9:35	Bert Still	Multi-lab	ECP Application Overview and Criteria	0:15
	9:50	BREAK			0:15
		NDA sessions	**These sessions require invidividuals or their institutions to be covered under proper NDA**		
	10:05	Intel NDA Session	Intel		1:05
	11:10	BREAK			0:10
	11:20	NVIDIA NDA Session	NVIDIA		1:05
	12:25	**LUNCH (on your own)**			1:20
Session chair:		**Apps / optimizations / algorithms**		**Application/algoroithm and/or platform-specific optimizations**	
Rebecca Hartman-Baker	13:45	Jae-Seung Yeom	LLNL	Data-dependent Performance Modeling of Linear Solvers for Sparse Matrices	0:15
	14:00	Charles Ferenbaugh	LANL	Coarse vs. fine-level threading in the PENNANT mini-app	0:15
	14:15	Scott Parker	ANL	Performance Optimization and Portability of the Nekbone Mini-App	0:15
	14:30	Kris Garrett	LANL	A first look at optimizing performance on the KNL	0:15
	14:45	Vitali Morozov	ANL	Portability of HACC - a highly tuned cosmology application	0:15
	15:00	BREAK			0:15
	15:15	Kristopher Keipert	ANL	Experiences and challenges while modernizing GAMESS for Theta and Aurora	0:15
	15:30	Steve Rennich	NVIDIA	GPU Performance Optimization of the Sweep Operation in Kripke	0:15
	15:45	Balint Joo	JLab/ANL/LBL	Experiences and Challenges for Performance Portability in Lattice QCD	0:15
	16:00	Alvaro Vazquez-Mayagoitia	ANL	Many-core and GPU developments in the parallel ELectronic Structure Infrastructure library (ELSI)	0:10
	16:10	BREAK			0:30
Session Chair:		**Performance Portable Abstractions**		**General abstractions suitable for managing portability in multiple applications**	
Hai Ah Nam	16:40	Tan Nguyen	LBL	Portable Data Locality Management with High-Level Programming Abstractions	0:15
	16:55	Jeff Vetter	ORNL	Understanding Portability of a High-Level Programming Model on Diverse HPC Architectures	0:20
	17:15	Christian Trott	SNL	Kokkos - Performance Portability Today	0:20
	17:35	Rich Hornung	LLNL	The RAJA Encapsulation Model for Architecture Portability	0:20
	17:55	Arpith Jacob	IBM	Towards Performance Portable GPU Programming with RAJA	0:20
	18:15	ADJOURN (dinner on your own)			

Day Two

	7:30	Coffee/Light Breakfast		Mingle	0:45
	8:15	Opening Remarks		Welcome, recap of day 1, overview of day 2, ….	0:05
Session Chair:		**Managing the Memory Hierachy**		**Abstractions/techniques for managing data motion between standard DRAM and HBM/Device memory**	
Katherine Riley	8:20	David Poliakoff	LLNL	Copy Hiding Application Interface (CHAI): Hiding Data Motion for Performance Portability	0:10

	8:30	Nikolai Sakharnykh	NVIDIA	Harnessing Performance of Geometric Multi-Grid Methods by using LOC and TOC architectures	0:15
	8:45	Fabian Delalondre	ANL	Leveraging heterogeneous systems and deep memory hierarchies for brain tissue modeling	0:20
	9:05	Luiz DeRose	Cray	Cray's Prog. Env. for Portable Performance and Programmability on Systems with High-Bandwidth Memory	0:15
	9:20	Ian Karlin	Multi-lab	Quad Lab Proposal of Fundamental Cross Architecture Multi-Level Memory Support	0:20
		Application Experience with Performance Portable Abstractions			
Session Chair:	9:40	Changhoan Kim	IBM	An abstraction for unstructured mesh problems	0:15
Tjerk Straatsma	9:55	Adam Kunen	LLNL	Nested Loop RAJA for Performance Portability	0:15
	10:10	Stan Moore	SNL	Obtaining Threading Performance Portability in SPARTA using Kokkos	0:15
	10:25	BREAK			0:30
	10:55	David Beckingsale	LLNL	Lightweight Models for Dynamically Tuning Data-Dependent Code	0:10
	11:05	Geoff Womeldorff	LANL	Kokkos and Legion Implementations of the SNAP Proxy Application	0:10
	11:15	Ryan Bleile	LLNL	Investigation of Portable Event-Based Monte Carlo Transport	0:15
	11:30	Matt Martineau	UK	Investigating the performance portability capabilities of OpenMP 4, Kokkos and Raja	0:20
	11:50	Leopold Grinberg	IBM	Performance portable single source-code implementation of sparse linear algebra operations on CPUs and GPUs	0:15
	12:05	Slaven Peles	LLNL	Investigating interoperability and performance portability of select LLNL numerical libraries	0:20
	12:25	LUNCH (provided / breakout topics by table)			0:55
Breakout #1	13:20	**BREAKOUT SESSION #1 (Managing the Memory Hierarchy / Performance Portable Abstractions)**			1:30
		Breakout Leads:	Doug Doerfler, Bronson Messer (mem hierarchy)		
			Brian Friesen, Jeff Vetter (PP abstractions)		
	14:50	BREAK			0:15
Session Chair:		**OpenMP**	**Experience with OpenMP and recommendations on guiding future standards**		
Hai Ah Nam	15:05	John Pennycook	Intel	Performance Portability of Kernel-based Abstractions	0:20
	15:25	John Pennycook	Intel	Generalizing a DSL for Structured Dependency (Stencil-like) Codes to OpenMP Loops	0:20
	15:45	John Levesque	Cray	How we can get Hybrid OpenMP/MPI to out perform All-MPI	0:20
	16:05	Carlo Bertolli	IBM	Performance Portability with OpenMP on Nvidia GPUs	0:20
	16:25	Jeff Larkin	NVIDIA	Performance Portability Through Descriptive Parallelism	0:20
	16:45	BREAK			0:30
	17:15	David Appelhans	IBM	Performance Portability Experience with LLVM, OpenMP 4, and Kripke	0:15
	17:30	Alexandre Eichenberger	IBM	OpenMP Specifications for Portability	0:15
	17:45	Oscar Hernandez	ORNL	Experiences with High-Level Programming Directives for Porting SPEC ACCEL on multiple architectures	0:15
	18:00	Tom Scogland	LLNL	Performance Portability with OpenMP: Experiences with 4.5 and Looking Toward 5.0	0:20
	18:20	Adjourn (dinner on your own)			
	19:30 - 22:00	Intel NDA Session		Optional set of evening talks on Intel NDA material for interested attendees	

Day Three

	7:30	Coffee/Light Breakfast		Mingle	0:45
	8:15	Recap of breakout #1		Each of four groups to present 8-10 minute summary	0:45
Session Chair:		**Tools / Compilers**	**Tools for performance portability and analysis**		
Hai Ah Nam	9:00	Jeanine Cook	SNL	The Importability of Performance Tools	0:10
	9:10	Juan Gonzalez Garcia	IBM	Next-gen profiling-infrastructure for supercomputers based on hybrid nodes	0:10
	9:20	Ignacio Laguna	LLNL	STATuner: Tuning CUDA Kernels via Compiler Analysis and Machine Learning	0:15
	9:35	Si Hammond	SNL	Profiling Interfaces for Parallel C++ Abstractions - KokkosP	0:15
	9:50	Protonu Basu	LBL	Leveraging Compiler-Based Tools for Performance-Portability	0:20
	10:10	Heidi Poxon	Cray	Adding Parallelism to HPC Applications using Reveal	0:15

	10:25	BREAK			0:30
Session Chair:		**IO / Burst Buffers**	**The I/O bottleneck and use of burst buffers**		
Mike Glass	10:55	Mark Miller	LLNL	Probing Portable Performance of Parallel I/O Paradigms using MACSio	0:15
	11:10	Andrey Ovsyannikov	LBL	ChomboCrunch and VisIt for carbon sequestration and in-transit data analysis using burst buffers	0:20
	11:30	Kathryn Mohror	LLNL	Performance Portability for Burst Buffers with the Scalable Checkpoint / Restart Library (SCR)	0:20
Session Chair:		**Domain Specific Languages**		**Use of DSL's for performance portability**	
Mike Glass	11:50	David Richards	LLNL	Portable Performance in Real Applications using Generated Code	0:15
	12:05	Brian Van Straalen	LBL	AMRStencil: An Embedded DSL for Expressing Structured Adaptive Mesh Refinement Algorithms	0:15
	12:20	LUNCH (provided / breakout topics by table)			0:55
Breakout #2	13:15	**BREAKOUT SESSION #2 (OpenMP Futures, Tools/Compiler/System Requirements)**			1:30
		Breakout Leads:	Sriram Swaminarayan, XXX (OpenMP Futures) Kathleen Shoga, TBD (Tools/Compilers/System)		
	14:45	BREAK		Scribes for breakout given time to collect notes	0:30
	15:15	**Recap of breakout #2**		Each of four groups to present 8-10 minute summary	0:45
		Wrapup discussions			
	16:00	Vendor Q&A / Panel		Vendor reps to discuss challenges and answer Q&A	0:40
	16:40	Wrapup / next-steps / takeaways		Capture followup goals, decide on subsequent meetings and potential topics	0:20
	17:00	ADJOURN			1:00
	18:00	DINNER			